arXiv:1402.6361v1 [math.OC] 25 Feb 2014

Oracle-Based Robust Optimization via Online Learning

Aharon Ben-Tal
Technion
abental@ie.technion.ac.il

Elad Hazan
Technion
ehazan@ie.technion.ac.il

Tomer Koren
Technion
tomerk@technion.ac.il

Shie Mannor
Technion
shie@ee.technion.ac.il

February 27, 2014

Abstract

Robust optimization is a common framework in optimization under uncertainty when the problem parameters are not known, but it is rather known that the parameters belong to some given uncertainty set. In the robust optimization framework the problem solved is a min-max problem where a solution is judged according to its performance on the worst possible realization of the parameters. In many cases, a straightforward solution of the robust optimization problem of a certain type requires solving an optimization problem of a more complicated type, and in some cases even NP-hard. For example, solving a robust conic quadratic program, such as those arising in robust SVM, ellipsoidal uncertainty leads in general to a semidefinite program. In this paper we develop a method for approximately solving a robust optimization problem using tools from online convex optimization, where in every stage a standard (non-robust) optimization program is solved. Our algorithms find an approximate robust solution using a number of calls to an oracle that solves the original (non-robust) problem that is inversely proportional to the square of the target accuracy.

1 Introduction

The Robust Optimization (RO; see Ben-Tal and Nemirovski (2002); Ben-Tal et al. (2009); Bertsimas et al. (2011)) framework addresses a fundamental problem of many convex optimization problems: slight inaccuracies in data give rise to significant fluctuations in the solution. While there are different approaches to handle uncertainty in the parameters of an optimization problem, the RO approach choose a solution that performs best against the worst possible parameter. When the objective function is convex in the parameters, and concave in the uncertainty, and when the uncertainty set is convex the overall optimization problem is convex.

Despite its theoretical and empirical success, a significant hinderance of adopting RO to large scale problems is the increased computational complexity. In particular, robust counterpart of an optimization problem is often more difficult, albeit usually convex, mathematical problems. For example, the robust counterpart of conic quadratic programming with ellipsoidal uncertainty constraints becomes a semi-definite program, for which we currently have significantly slower solvers.

RO has recently gained traction as a tool for analyzing machine learning algorithms and for devising new ones. In a sequence of papers, Xu, Caramanis and Mannor show that several standard machine learning algorithms such as Lasso and norm regularized support vector machines have a RO interpretation (Xu et al., 2009, 2010). Beyond these works, robustness is a desired property for many learning algorithms. Indeed, making standard algorithms robust to outliers or to perturbation in the data has been proposed in several works; see Lanckriet et al. (2003); Bhattacharyya et al. (2004a,b); Shivaswamy et al. (2006); Trafalis and Gilbert (2007). However in these cases, the problem eventually solved is more complicated than the original problem. For example, in Trafalis and Gilbert (2007) the original problem is a standard support vector machine, but when robustifying it to input uncertainty, one has to solve a second-order conic program (in the non-separable case). Another example is Shivaswamy et al. (2006) where the uncertainty is a probability distribution over inputs. In that case, the original SVM becomes a second-order conic program as well.

The following question arrises: **can we (approximately) solve a robust counterpart of a given optimization problem using only an algorithm for the original optimization formulation**? In this paper we answer this question on the affirmative: we give two meta-algorithms that receive as input an oracle to the original mathematical problem and approximates the robust counterpart by invoking the oracle a polynomial number of times. In both approaches, the number of iterations to obtain an approximate robust solution is a function of the approximation guarantee and the complexity of the uncertainty set, and does not directly depend on the dimension of the problem. Our methods differ on the assumptions regarding the uncertainty set and the dependence of the constraints on the uncertainty. The first method allows any concave function of the noise terms but is limited to convex uncertainty sets. The second method allows arbitrary uncertainty sets as long as a "pessimization oracle" (as termed by Mutapcic and Boyd 2009) exists— an oracle that finds the worst case-noise for a given feasible solution. Our methods are formally described as template, or meta-algorithms, and are general enough to be applied even if the robust counterpart is NP-hard[1].

Our approach for achieving efficient oracle-based RO is to reduce the robust formulation to a zero-sum game, which we solve by a primal-dual technique based on tools from online learning. Such primal-dual methods originated from the study of approximation algorithms for linear programs (Plotkin et al., 1995) and were recently proved invaluable in understanding Lagrangian relaxation methods (see e.g. Arora et al. 2012) and in sublinear-time optimization techniques (Clarkson et al., 2012; Hazan et al., 2011; Garber and Hazan,

[1]Recall that we are only providing an approximate solution, and thus our algorithms formally constitute a "polynomial time approximation scheme" (PTAS).

2011). We show how to apply this methodology to oracle-based RO. Along the way, we contribute some extensions to the existing online learning literature itself, notably giving a new Follow-the-Perturbed-Leader algorithm for regret minimization that works with (additive) approximate linear oracles.

Finally, we demonstrate examples and applications of our methods to various RO formulation including linear, semidefinite and quadratic programs. The latter application builds on recently developed efficient linear-time algorithms for the trust region problem (Hazan and Koren, 2014).

Related work. Robust optimization is by now a field of study by itself, and the reader is referred to Ben-Tal et al. (2009); Bertsimas et al. (2011) for further information and references. The computational bottleneck associated with robust optimization was addressed in several papers. Calafiore and Campi (2004) propose to sample constraints from the uncertainty set, and obtain an "almost-robust" solution with high probability with enough samples. The main problem with their approach is that the number of samples can become large for a high-dimensional problem.

For certain types of discrete robust optimization problems, Bertsimas and Sim (2003) propose solving the robust version of the problem via n (the dimension) solutions of the original problem. Mutapcic and Boyd (2009) give an iterative cutting plane procedure for attaining essentially the same goal as us, and demonstrate impressive practical performance. However, the overall running time for their method can be exponential in the dimension.

Oraganization. The rest of the paper is organized as follows. In Section 2 we present the model and set the notations for the rest of the paper. In Section 3.1 we describe the simpler of our two meta-algorithms: a meta algorithm for approximately solving RO problems that employs dual subgradient steps, under the assumption that the robust problem is convex with respect to the noise variables. In Section 3.2 we remove the latter convexity assumption and only assume that the problem of finding the worst-case noise assignment can be solved by invoking a "pessimization oracle". This approach is more general than the subgradient-based method and we exhibit perhaps our strongest example of solving robust quadratic programs using this technique in Section 4. Section 4 also contains examples of application of our technique for robust linear and semi-definite programming. We conclude in Section 5.

2 Preliminaries

We start this section with the standard formulation of RO. We then recall some basic results from online learning.

2.1 Robust Optimization

Consider a convex mathematical program in the following general formulation:

$$\begin{array}{ll} \text{minimize} & f_0(x) \\ \text{subject to} & f_i(x, u_i) \leq 0, \qquad i = 1, \ldots, m \ , \\ & x \in \mathcal{D} \ . \end{array} \tag{1}$$

Here $f_0, f_1, \ldots, f_m$ are convex functions, $\mathcal{D} \subseteq \mathbb{R}^n$ is a convex set in Euclidean space, and $\boldsymbol{u} = (u_1, \ldots, u_m)$ is a fixed parameter vector. The *robust counterpart* of this formulation is given by

$$\begin{array}{ll} \text{minimize} & f_0(x) \\ \text{subject to} & f_i(x, u_i) \leq 0, \quad \forall \ u_i \in \mathcal{U}, \ \ i = 1, \ldots, m \ , \\ & x \in \mathcal{D} \ , \end{array} \tag{2}$$

where the parameter vector $\boldsymbol{u}$ is constrained to be in a set $\boldsymbol{\mathcal{U}} = \mathcal{U} \times \cdots \times \mathcal{U} = \mathcal{U}^m$ called the *uncertainty set*. It is without loss of generality to assume that the uncertainty set has this specific form of a cartesian product, see e.g. Ben-Tal and Nemirovski (2002). Here we also assume that the uncertainty set is symmetric (that is, its projection onto each dimension is the same set). This assumption is only made for simplifying notations and can be relaxed easily.

The following observation is standard: we can reduce the above formulation to a feasibility problem via a binary search over the optimal value of (1), replacing the objective with the constraint $f_0(x) - t \leq 0$ with t being our current guess of the optimal value (of course, assuming the range of feasible values is known a-priory). For ease of notation, we rename f_0 by shifting it by α, and can write the first constraint as simply $f_0(x, u) \leq 0$. With these observations, we can reduce the robust counterpart to the feasibility problem

$$\exists ?\ x \in \mathcal{D} \quad : \quad f_i(x, u_i) \leq 0 \ , \quad \forall \ u_i \in \mathcal{U}, \ \ i = 1, \ldots, m \ . \tag{3}$$

We say that $x \in \mathcal{D}$ is an *ϵ-approximate solution* to this problem if x meets each constraint up to ϵ, that is, if it satisfies $f_i(x, u_i) \leq \epsilon$ for all $u_i \in \mathcal{U}$ $(i = 1, \ldots, m)$.

2.2 Online Convex Optimization and Regret minimization

Our derivations below use tools from online learning, namely algorithms for minimizing regret in the general prediction framework of Online Convex Optimization (OCO). In OCO[2], the online predictor iteratively produces a decision $x_t \in \mathcal{K}$ from a convex decision set $\mathcal{K} \subseteq \mathbb{R}^n$. After a decision is generated, a concave reward function f_t is revealed, and the decision maker suffers a loss of $f_t(x_t) \in \mathbb{R}$. The standard performance metric in online learning is called *regret*, given by

$$R_T \ = \ \max_{x^* \in \mathcal{K}} \sum_{t=1}^{T} f_t(x^*) \ - \ \sum_{t=1}^{T} f_t(x_t) \ .$$

[2]Here we present OCO as the problem of online maximization of *concave reward* functions rather than online minimization of *convex cost* functions. While the latter is more common, both formulations are equivalent.

The reward function is not known to the decision maker before selecting x_t and it is, in principal, arbitrary and even possibly chosen by an adversary. We henceforth make crucial use of this robustness against adversarial choice of reward functions: the reward functions we shall use will be chosen by a dual optimization problem, thereby directing the entire algorithm towards a correct solution. We refer the reader to (Cesa-Bianchi and Lugosi, 2006; Hazan, 2011; Shalev-Shwartz, 2012) for more details on online learning and online convex optimization.

Two regret-minimization algorithms that we shall use henceforth (at least in spirit) are Online Gradient Descent (Zinkevich, 2003) and Follow the Perturbed Leader (Kalai and Vempala, 2005).

Online Gradient Descent (OGD). In OGD, the decision maker predicts according to the rule

$$x_{t+1} \;\leftarrow\; \Pi_{\mathcal{K}}\big[x_t + \eta \nabla f_t(x_t)\big] \;, \tag{4}$$

where $\Pi_{\mathcal{K}}(x) = \min_{y\in\mathcal{K}} \|x - y\|$ is the Euclidean projection operator onto the set $\mathcal{K}$. Hence, the OGD algorithm takes a projected step in the direction of the gradient of the current reward function. Even thought the next reward function can be arbitrary, it can be shown that this algorithm achieves a sublinear regret.

Lemma 1 (Zinkevich 2003)**.** *For any sequence of reward functions, let $\{x_t\}$ be the sequence generated by* (4)*. Then, setting $\eta = D/G\sqrt{T}$ we obtain*

$$\max_{x^*\in\mathcal{K}} \sum_t f_t(x^*) \;-\; \sum_{t=1}^{T} f_t(x_t) \;\le\; GD\sqrt{T} \;,$$

where $G \ge \max_{t\in[T]} \|f_t(x_t)\|$ is an upper bound on the ℓ_2 norm of the gradients of the reward functions, and $D \ge \max_{x,y\in\mathcal{K}} \|x - y\|$ is an upper bound on the ℓ_2 diameter of $\mathcal{K}$.

Follow the Perturbed Leader (FPL). The FPL algorithm works in a similar setting as OGD, but with two crucial differences:

1. The set $\mathcal{K}$ does not need to be convex. This is a significant advantage of the FPL approach, which we make use of in our application to robust quadratic programming (see Section 4.2).

2. FPL assumes that the reward functions $f_1, \ldots, f_T$ are linear, i.e. $f_t(x) = f_t \cdot x$ with $f_t \in \mathbb{R}^n$.

Kalai and Vempala (2005) suggest the following method for online decision making that relies on a linear optimization procedure M over the set $\mathcal{K}$ that computes $M(f) = \arg\max_{x\in\mathcal{K}} f \cdot x$ for all $f \in \mathbb{R}^n$. FPL chooses x_t by first drawing a perturbation $p_t \in [0, \frac{1}{\eta}]^n$ uniformly at random, and computing:

$$x_{t+1} \;=\; M\left(\sum\nolimits_{\tau=1}^{t} f_\tau + p_t\right) \;. \tag{5}$$

The regret of this algorithm is bounded as follows.

Lemma 2 (Kalai and Vempala 2005)**.** *For any sequence of reward vectors $f_1, \ldots, f_T$, let $x_1, \ldots, x_T$ be the sequence of decisions generated by (5) with parameter $\eta = \sqrt{D/RAT}$. Then*

$$\max_{x^* \in \mathcal{K}} \sum_{t=1}^{T} f_t \cdot x^* \; - \; \mathbf{E}\left[\sum_{t=1}^{T} f_t \cdot x_t\right] \; \leq \; 2\sqrt{DRAT} \; ,$$

where $R \geq \max_{t,x} |f_t \cdot x|$ is an upper bound on the magnitude of the rewards, $A \geq \max_t \|f_t\|_1$ is an upper bound on the ℓ_1 norm of the reward vectors, and $D \geq \max_{x,y \in \mathcal{K}} \|x - y\|_1$ is an upper bound on the ℓ_1 diameter of $\mathcal{K}$.

For our purposes, and in order to be able to work with an approximate optimization oracle to the original mathematical program, we need to adapt the original FPL algorithm to work with noisy oracles. This adaptation is made precise in Section 3.2.

3 Oracle-Based Robust Optimization

In this section we formally state and prove our first (and simpler) result: an oracle-based approximate robust optimization algorithm that is based on subgradient descent.

Throughout the section we assume the availability of an optimization oracle for the original optimization problem of the form given in Figure 1, which we denote by $\mathcal{O}_\epsilon$. Such an optimization oracle approximately solves formulation (3) for any *fixed* noise $\boldsymbol{u} \in \boldsymbol{\mathcal{U}}$, in the sense that it either returns an ϵ-feasible solution x (that meets each constraint up to ϵ) or correctly declares that the problem is infeasible.

Input: noise vector $\boldsymbol{u} = (u_1, \ldots, u_m) \in \boldsymbol{\mathcal{U}}$
Output: vector $x \in \mathcal{D}$ that satisfies

$$f_i(x, u_i) \leq \epsilon \; , \quad i = 1, \ldots, m$$

or "INFEASIBLE" if there does not exist a vector $x \in \mathcal{D}$ for which

$$f_i(x, u_i) \leq 0 \; , \quad i = 1, \ldots, m$$

Figure 1: An ϵ-approximate optimization oracle for the original optimization problem.

3.1 Dual-Subgradient Meta-Algorithm

In this section we assume that for all $i = 1, \ldots, m$:

1. For all $x \in \mathcal{D}$, the function $f_i(x, u)$ is *concave* in u;
2. The set $\mathcal{U}$ is *convex*.

Under these assumptions, the robust formulation is in fact a convex-concave saddle-point problem that can be solved in polynomial time using interior-point methods. However, recall that our goal is to solve the robust problem by invoking a solver of the original (non-robust) optimization problem.

In the setting of this section, we shall make use of the following definitions. Let D be an upper bound over the ℓ_2 diameter of $\mathcal{U}$, that is $D \geq \max_{u,v \in \mathcal{U}} \|u - v\|_2$. Let G be a constant such that $\|\nabla_u f_i(x,u)\|_2 \leq G$ for all $x \in \mathcal{D}$ and $u \in \mathcal{U}$.

Algorithm 1 DUAL-SUBGRADIENT ROBUST OPTIMIZATION

input: target accuracy $\epsilon > 0$, parameters D, G
output: 2ϵ-approximate solution to (3), or "INFEASIBLE"

set $T = \lceil G^2 D^2 / \epsilon^2 \rceil$ and $\eta = D/G\sqrt{T}$
initialize $(u_1^0, \ldots, u_m^0) \in \boldsymbol{\mathcal{U}}$ arbitrarily
for $t = 1$ to T **do**
  for $i = 1$ to m **do**
    update $u_i^t \leftarrow \Pi_{\mathcal{U}}\big[\, u_i^{t-1} + \eta \nabla_u f_i(x^{t-1}, u_i^{t-1}) \,\big]$
  end for
  set $x^t \leftarrow \mathcal{O}_\epsilon(u_1^t, \ldots, u_m^t)$
  if oracle declared infeasibility **then return** "INFEASIBLE"
end for
return $\bar{x} = \frac{1}{T} \sum_{t=1}^{T} x^t$

With the above assumptions and definitions, we can now present an oracle-based robust optimization algorithm, given in Algorithm 1. The algorithm is comprised of primal-dual iterations, where the dual part of the algorithm updates the noise terms according to the current primal solution, via a low-regret update. For this algorithm, we prove:

Theorem 3. *Algorithm 1 either returns an 2ϵ-approximate solution to the robust program (3) or correctly concludes that it is infeasible. The algorithm terminates after at most $O(G^2 D^2 / \epsilon^2)$ calls to the oracle $\mathcal{O}_\epsilon$.*

Proof. First, suppose that the algorithm returns "INFEASIBLE". By the definition of the oracle $\mathcal{O}_\epsilon$, this happens if for some $t \in [T]$, there does not exists $x \in \mathbb{R}^n$ such that

$$f_i(x, u_i^t) \leq 0 \ , \quad i = 1, \ldots, m \ .$$

This implies that the robust counterpart (3) cannot be feasible, as there exists an admissible perturbation that makes the original problem infeasible.

Next, suppose that a solution $\bar{x}$ is returned. The premise of the oracle implies that $f_i(x^t, u_i^t) \leq \epsilon$ for all $t \in [T]$ and $i \in [m]$ (otherwise, the algorithm would have returned "INFEASIBLE"), whence

$$\forall \ i \in [m] \ , \qquad \frac{1}{T} \sum_{t=1}^{T} f_i(x^t, u_i^t) \ \leq \ \epsilon \ . \tag{6}$$

On the other hand, from the regret guarantee of the Online Gradient Descent algorithm we have

$$\forall\ i \in [m]\ , \qquad \max_{u_i \in \mathcal{U}} \frac{1}{T}\sum_{t=1}^{T} f_i(x^t, u_i)\ -\ \frac{1}{T}\sum_{t=1}^{T} f_i(x^t, u_i^t)\ \leq\ \frac{GD}{\sqrt{T}}\ \leq\ \epsilon\ . \tag{7}$$

Combining (10) and (12), we conclude that for all $i \in [m]$,

$$\epsilon\ \geq\ \frac{1}{T}\sum_{t=1}^{T} f_i(x^t, u_i^t)\ \geq\ \max_{u_i \in \mathcal{U}} \frac{1}{T}\sum_{t=1}^{T} f_i(x^t, u_i) - \epsilon\ \geq\ \max_{u_i \in \mathcal{U}} f_i(\bar{x}, u_i) - \epsilon\ ,$$

where the final inequality follows from the convexity of the functions f_i with respect to x. Hence, for every $i \in [m]$ we have

$$f_i(\bar{x}, u_i)\ \leq\ 2\epsilon, \qquad \forall\ u_i \in \mathcal{U}\ ,$$

implying that $\bar{x}$ is an 2ϵ-approximate robust solution. □

3.2 Dual-Perturbation Meta-Algorithm

We now give our more general and intricate oracle-based approximation algorithm for RO. In contrast to the previous simple subgradient-based method, in this section we do not need the uncertainty structure to be convex. Instead, in addition to an oracle to solve the original mathematical program, we also assume the existence of an efficient "pessimization oracle" (as termed by Mutapcic and Boyd 2009), namely an oracle that approximates the worst-case noise for any given feasible solution x. Formally, assume that for all $i = 1, 2, \ldots, m$ the following hold:

1. For all $x \in \mathcal{D}$, the function $f_i(x, u)$ is *linear* in u, i.e. can be written as $f_i(x, u) = \boldsymbol{g}_i(x) \cdot u + h_i(x)$ for some functions $\boldsymbol{g}_i : \mathbb{R}^n \mapsto \mathbb{R}^d$ and $h_i : \mathbb{R}^n \mapsto \mathbb{R}$;

2. There exists a linear optimization procedure M_ϵ that given a vector $g \in \mathbb{R}^d$, computes a vector $M_\epsilon(g) \in \mathbb{R}^d$ such that $g \cdot M_\epsilon(g) \geq \arg\max_{u \in \mathcal{U}} g \cdot u - \epsilon$.

On the surface, the linearity assumption seems very strong. However, note that *we do not assume the convexity of the set* $\mathcal{U}$. This means that the dual subproblem (that amounts to finding the worst-case noise for a given x) is not necessarily a convex program. Nevertheless, our approach can still approximate the robust formulation as long as a procedure M_ϵ is available.

In the rest of the section we use the following notations. Let D be an upper bound over the ℓ_1 diameter of $\mathcal{U}$, that is $D \geq \max_{u,v \in \mathcal{U}} \|u - v\|_1$. Let F and G be constants such that $|f_i(x, u)| \leq F$ and $\|\boldsymbol{g}_i(x)\|_1 \leq G$ for all $x \in \mathcal{D}$ and $u \in \mathcal{U}$.

We can now present our second oracle-based meta-algorithm, described in Algorithm 2. Similarly to our dual-subgradient method, the algorithm is based on primal-dual iterations. However, in the dual part we now rely on the approximate pessimization oracle for updating the noise terms. This algorithm provides the following convergence guarantee.

Algorithm 2 DUAL-PERTURBATION ROBUST OPTIMIZATION

input: target accuracy and confidence $\epsilon, \delta > 0$, parameters D, F, G
output: 4ϵ-approximate solution to (3), or "INFEASIBLE"

set $T = \lceil \max\{DG, F\} \cdot \frac{16F}{\epsilon^2} \log \frac{m}{\delta} \rceil$ and $\eta = \sqrt{D/FGT}$
for $t = 1$ to T **do**
  for $i = 1$ to m **do**
    choose $p_t \in [0, \frac{1}{\eta}]^n$ uniformly at random
    compute $u_i^t \leftarrow M_\epsilon\big(\sum_{\tau=1}^{t} \boldsymbol{g}_i(x^\tau) + p_t \big)$
  end for
  set $x^t \leftarrow \mathcal{O}_\epsilon(u_1^t, \ldots, u_m^t)$
  if oracle declared infeasibility **then return** "INFEASIBLE"
end for
return $\bar{x} = \frac{1}{T} \sum_{t=1}^{T} x^t$

Theorem 4. *With probability at least $1 - \delta$, Algorithm 2 either returns an 4ϵ-approximate solution to the robust program (3) or correctly concludes that it is infeasible. The algorithm terminates after at most $T = \tilde{O}((DG + F) \cdot F/\epsilon^2)$ calls to the oracle $\mathcal{O}_\epsilon$.*

We begin by analyzing the dual part of the algorithm, namely, the rule by which the variables u_i^t are updated. While this rule is essentially an FPL-like update, we cannot apply Lemma 2 directly for two crucial reasons. First, the update uses an approximate linear optimization procedure instead of an exact one as required by FPL. Second, the reward vectors $\boldsymbol{g}_i(x^1), \ldots, \boldsymbol{g}_i(x^T)$ being observed by the dual algorithm are random variables that depend on its internal randomization (i.e., on the random variables $p_1, \ldots, p_T$). Nevertheless, by analyzing a noisy version of the FPL algorithm (in Section 3.3 below) we can prove the following bound.

Lemma 5. *For each $i = 1, \ldots, m$, with probability at least $1 - \delta$ we have that*

$$\max_{u_i \in \mathcal{U}} \sum_{t=1}^{T} \boldsymbol{g}_i(x^t) \cdot u_i \;-\; \sum_{t=1}^{T} \boldsymbol{g}_i(x^t) \cdot u_i^t \;\le\; 2\sqrt{DFGT} + 2F\sqrt{T \log \tfrac{1}{\delta}} + 2\epsilon T \;.$$

Proof. Fix some $i \in [m]$. Note that the distribution from which the dual algorithm draws u_i^t is a deterministic function of the primal variables $x^1, \ldots, x^{t-1}$. Hence, we can apply Lemma 4.1 of Cesa-Bianchi and Lugosi (2006), together with the regret bound of Theorem 6 (see Section 3.3 below), and obtain that

$$\max_{u_i \in \mathcal{U}} \sum_{t=1}^{T} \boldsymbol{g}_i(x^t) \cdot u_i \;-\; \sum_{t=1}^{T} \mathbf{E}_t[\boldsymbol{g}_i(x^t) \cdot u_i^t] \;\le\; 2\sqrt{DFGT} + 2\epsilon T \;, \tag{8}$$

where $\mathbf{E}_t[\cdot]$ denotes the expectation conditioned on $p_1, \ldots, p_{t-1}$. Next, note that the random variables $Z_t = \boldsymbol{g}_i(x^t) \cdot u_i^t - \mathbf{E}_t[\boldsymbol{g}_i(x^t) \cdot u_i^t]$ for $t = 1, \ldots, T$ form a martingale differences

sequence with respect to $p_1, \ldots, p_T$, and

$$|Z_t| \;\leq\; \left|\boldsymbol{g}_i(x^t)\cdot u_i^t\right| + \mathbf{E}_t[\left|\boldsymbol{g}_i(x^t)\cdot u_i^t\right|] \;\leq\; 2F \;.$$

Hence, by Azuma's inequality (see e.g., Lemma A.7 in Cesa-Bianchi and Lugosi 2006), with probability at least $1-\delta$,

$$\sum_{t=1}^{T} \mathbf{E}_t[\boldsymbol{g}_i(x^t)\cdot u_i^t] \;-\; \sum_{t=1}^{T} \boldsymbol{g}_i(x^t)\cdot u_i^t \;\leq\; 2F\sqrt{T\log\tfrac{1}{\delta}} \;. \tag{9}$$

Summing inequalities (8) and (9), we obtain the lemma. □

Equipped with the above lemma, we can now prove Theorem 4.

Proof of Theorem 4. First, suppose that the algorithm returns "INFEASIBLE". By the definition of the oracle $\mathcal{O}_\epsilon$, this happens if for some $t \in [T]$, there does not exists $x \in \mathbb{R}^n$ such that

$$f_i(x, u_i^t) \leq 0 \;, \quad i = 1, \ldots, m \;.$$

This implies that the robust counterpart (3) cannot be feasible.

Next, suppose that a solution $\bar{x}$ is returned (note that $\bar{x}$ must lie in the set $\mathcal{D}$ as we assume that $\mathcal{D}$ is convex). This ensures that $f_i(x^t, u_i^t) \leq \epsilon$ for all $t \in [T]$ and $i \in [m]$ (otherwise, the algorithm would have returned "INFEASIBLE"), whence

$$\forall\; i \in [m] \;, \qquad \frac{1}{T}\sum_{t=1}^{T} f_i(x^t, u_i^t) \;\leq\; \epsilon \;. \tag{10}$$

On the other hand, Lemma 5 implies that for each $i \in [m]$ we have

$$\max_{u_i \in \mathcal{U}} \frac{1}{T}\sum_{t=1}^{T} \boldsymbol{g}_i(x^t)\cdot u_i \;-\; \frac{1}{T}\sum_{t=1}^{T} \boldsymbol{g}_i(x^t)\cdot u_i^t \;\leq\; 2\sqrt{\frac{DFG}{T}} + 2F\sqrt{\frac{\log\frac{m}{\delta}}{T}} + 2\epsilon$$

with probability at least $1-\delta/m$. Recalling that $f_i(x^t, u) = \boldsymbol{g}_i(x^t)\cdot u$ for all $u \in \mathcal{U}$ and applying a union bound, we obtain that with probability at least $1-\delta$,

$$\forall\; i \in [m], \quad \max_{u_i \in \mathcal{U}} \frac{1}{T}\sum_{t=1}^{T} f_i(x^t, u_i) \;-\; \frac{1}{T}\sum_{t=1}^{T} f_i(x^t, u_i^t) \;\leq\; 2\sqrt{\frac{DFG}{T}} + 2F\sqrt{\frac{\log\frac{m}{\delta}}{T}} + 2\epsilon \;. \tag{11}$$

Using our choice of T now gives that with probability at least $1-\delta$,

$$\forall\; i \in [m], \quad \max_{u_i \in \mathcal{U}} \frac{1}{T}\sum_{t=1}^{T} f_i(x^t, u_i) \;-\; \frac{1}{T}\sum_{t=1}^{T} f_i(x^t, u_i^t) \;\leq\; 3\epsilon \;. \tag{12}$$

Combining (10) and (12), we conclude that with probability at least $1-\delta$, for all $i \in [m]$,

$$\epsilon \;\geq\; \frac{1}{T}\sum_{t=1}^{T} f_i(x^t, u_i^t) \;\geq\; \max_{u_i \in \mathcal{U}} \frac{1}{T}\sum_{t=1}^{T} f_i(x^t, u_i) - 3\epsilon \;\geq\; \max_{u_i \in \mathcal{U}} f_i(\bar{x}, u_i) - 3\epsilon \;,$$

where the final inequality follows from the convexity of the functions f_i with respect to x. Hence, with probability at least $1-\delta$, for every $i \in [m]$ we have

$$f_i(\bar{x}, u_i) \;\leq\; 4\epsilon, \qquad \forall\; u_i \in \mathcal{U} \;,$$

implying that $\bar{x}$ is an 4ϵ-approximate robust solution. □

3.3 Follow the Approximate Perturbed Leader

As mentioned above, in our analysis we require a noisy version of the FPL algorithm, namely a variant capable of using an approximate linear optimization procedure over the decision domain rather than an exact one. Here we analyze such a variant and prove Theorem 6 being used in the proof of Lemma 5 above.

Assume we have a procedure M_ϵ for ϵ-approximating linear programs over a (not necessarily convex) domain $\mathcal{K}$, that is, for all $f \in \mathbb{R}^n$ the output of $M_\epsilon(f)$ satisfies

$$f \cdot M_\epsilon(f) \;\geq\; \max_{x \in \mathcal{K}} f \cdot x \;-\; \epsilon$$

for some constant $\epsilon > 0$. We analyze the following version of the FPL algorithm: at round t choose x_t by first choosing a perturbation $p_t \in [0, 1/\eta]^n$ uniformly at random, and computing:

$$x_{t+1} \;=\; M_\epsilon\left(\sum\nolimits_{\tau=1}^{t} f_t + p_t\right) \;. \tag{13}$$

We show that the error introduced by the noisy optimization procedure M_ϵ does not harm the regret too much. Formally, we prove:

Theorem 6. *For any sequence of reward vectors $f_1, \ldots, f_T$, let $x_1, \ldots, x_T$ be the sequence of decisions produced by (13) with parameter $\eta = \sqrt{D/RAT}$. Then*

$$\max_{x^* \in \mathcal{K}} \sum_{t=1}^{T} f_t \cdot x^* \;-\; \mathbf{E}\left[\sum_{t=1}^{T} f_t \cdot x_t\right] \;\leq\; 2\sqrt{DRAT} + 2\epsilon T \;,$$

where $R \geq \max_{t,x} |f_t \cdot x|$ is an upper bound on the magnitude of the rewards, $A \geq \max_t \|f_t\|_1$ is an upper bound on the ℓ_1 norm of the reward vectors, and $D \geq \max_{x,y \in \mathcal{K}} \|x - y\|_1$ is an upper bound on the ℓ_1 diameter of $\mathcal{K}$.

Throughout this section we use the notation $f_{1:t}$ as a shorthand for the sum $\sum_{\tau=1}^{t} f_\tau$. Following the analysis of Kalai and Vempala (2005), we first prove that being the approximate leader yields approximately zero regret.

Lemma 7. *For any sequence of vectors $f_1, \ldots, f_T$,*

$$\sum_{t=1}^{T} M_\epsilon(f_{1:t}) \cdot f_t \;\geq\; M_\epsilon(f_{1:T}) \cdot f_{1:T} - \epsilon T \ .$$

Proof. The proof is by induction on T. For $T = 1$ the claim is trivial. Next, assuming correctness for some value of T we have

$$\begin{aligned}
\sum_{t=1}^{T+1} M_\epsilon(f_{1:t}) \cdot f_t \;&\geq\; M_\epsilon(f_{1:T}) \cdot f_{1:T} + M_\epsilon(f_{1:T+1}) \cdot f_{T+1} - \epsilon T \\
&\geq\; M_\epsilon(f_{1:T+1}) \cdot f_{1:T} - \epsilon + M_\epsilon(f_{1:T+1}) \cdot f_{T+1} - \epsilon T \\
&=\; M_\epsilon(f_{1:T+1}) \cdot f_{1:T+1} - \epsilon(T+1) \ ,
\end{aligned}$$

which completes the proof. □

Next, we bound the regret of a hypothetical algorithm that on round t uses the unobserved function f_t for predicting x_t.

Lemma 8. *For any vectors $f_1, \ldots, f_T$ ($T \geq 2$) and $p \in [0, 1/\eta]^n$ it holds that*

$$\sum_{t=1}^{T} M_\epsilon(f_{1:t} + p) \cdot f_t \;\geq\; \max_{x \in \mathcal{K}} f_{1:t} \cdot x - \frac{D}{\eta} - 2\epsilon T \ ,$$

where $D \geq \max_{x,y \in \mathcal{K}} \|x - y\|_1$ is an upper bound on the ℓ_1 diameter of $\mathcal{K}$.

Proof. Imagine a fictitious round $t = 0$ in which a reward vector $f_0 = p$ is observed. Then, using Lemma 7 we can write

$$\begin{aligned}
\sum_{t=1}^{T} M_\epsilon(f_{1:t} + p) \cdot f_t \;&=\; \sum_{t=0}^{T} M_\epsilon(f_{0:t}) \cdot f_t - M_\epsilon(f_0) \cdot f_0 \\
&\geq\; M_\epsilon(f_{0:t}) \cdot f_{0:t} - M_\epsilon(f_0) \cdot f_0 - \epsilon T \ .
\end{aligned}$$

Using the guarantee of M_ϵ, we can bound the first term on the right hand side as

$$\begin{aligned}
M_\epsilon(f_{0:t}) \cdot f_{0:t} \;&\geq\; M_\epsilon(f_{1:t}) \cdot f_{0:t} - \epsilon \\
&=\; M_\epsilon(f_{1:t}) \cdot f_{1:t} + M_\epsilon(f_{1:t}) \cdot f_0 - \epsilon \\
&\geq\; \max_{x \in \mathcal{K}} f_{1:t} \cdot x + M_\epsilon(f_{1:t}) \cdot f_0 - 2\epsilon \ .
\end{aligned}$$

Putting things together, for $T \geq 2$ we have

$$\begin{aligned}
\sum_{t=1}^{T} M_\epsilon(f_{1:t} + p) \cdot f_t \;&\geq\; \max_{x \in \mathcal{K}} f_{1:t} \cdot x + (M_\epsilon(f_{1:t}) - M_\epsilon(f_0)) \cdot f_0 - (T+2)\epsilon \\
&\geq\; \max_{x \in \mathcal{K}} f_{1:t} \cdot x - \frac{D}{\eta} - 2\epsilon T \ ,
\end{aligned}$$

where the final inequality follows from Hölder's inequality, since $\|M_\epsilon(f_{1:t}) - M_\epsilon(f_0)\|_1 \leq D$ and $\|f_0\|_\infty = \|p\|_\infty \leq 1/\eta$. □

Our final lemma bounds the expected difference in quality between the prediction made by the hypothetical algorithm to the one made by the approximate FPL algorithm.

Lemma 9. *For any sequence $f_1, \ldots, f_t$ of reward vectors and p distributed uniformly in the cube $[0, 1/\eta]^n$ we have*

$$\mathbf{E}[M_\epsilon(f_{1:t-1} + p) \cdot f_t] \; - \; \mathbf{E}[M_\epsilon(f_{1:t} + p) \cdot f_t] \; \geq \; -\eta RA \; ,$$

where $R \geq \max_{t,x} |f_t \cdot x|$ and $A \geq \max_t \|f_t\|_1$.

Proof. Lemma 3.2 in Kalai and Vempala (2005) shows that the cubes $f_{1:t-1} + p$ and $f_{1:t} + p$ overlap in at least $1 - \eta|f_t| \geq 1 - \eta A$ fraction. On this intersection, the random variables $M_\epsilon(f_{1:t-1} + p) \cdot f_t$ and $M_\epsilon(f_{1:t} + p) \cdot f_t$ are identical. Otherwise, they can differ by at most R. This gives the claim. □

We can now prove our regret bound.

Proof of Theorem 6. Since we are bounding the expected regret, we can simply assume that $p_1 = \ldots = p_T = p$ with p uniformly distributed in the cube $[0, 1/\eta]^n$. Combining the above lemmas, we see that

$$\begin{aligned}
\mathbf{E}\left[\sum_{t=1}^{T} M_\epsilon(f_{1:t-1} + p_t) \cdot f_t\right] \; &= \; \mathbf{E}\left[\sum_{t=1}^{T} M_\epsilon(f_{1:t-1} + p) \cdot f_t\right] \\
&\geq \; \mathbf{E}\left[\sum_{t=1}^{T} M_\epsilon(f_{1:t} + p) \cdot f_t\right] - \eta RAT \\
&\geq \; \max_{x \in \mathcal{K}} f_{1:t} \cdot x - \frac{D}{\eta} - \eta RAT - 2\epsilon T \; .
\end{aligned}$$

The claimed regret bound now follows from our choice of $\eta = \sqrt{D/RAT}$. □

4 Examples and Applications

In this section we provide several examples for the applicability of our results. All the problems we consider are stated as feasibility problems. For concreteness, we focus on ellipsoidal uncertainty sets, being the most common model of data uncertainty.

4.1 Robust Linear Programming

A linear program (LP) in the standard form is given by

$$\begin{aligned}
\exists? \quad & x \in \mathbb{R}^n \\
\text{s.t.} \quad & a_i^\top x - b_i \leq 0 \; , \quad i = 1, \ldots, m \; ,
\end{aligned}$$

The robust counterpart of this optimization problem is a second-order conic program (SOCP) that can be solved efficiently, see e.g. Ben-Tal et al. (2009); Bertsimas et al. (2011). In many cases of interest there exist highly efficient solvers for the original LP problem, as in the important case of network flow problems where the special combinatorial structure allows for algorithms that are much more efficient than generic LP solvers. However, this combinatorial structure is lost for its corresponding robust network flow problem. Hence, solving the robust problem using an oracle-based approach might be favorable in these cases. For the same reason, our technique is relevant even in the case of polyhedral uncertainty, where the robust counterpart remains an LP but possibly without the special structure of the original formulation.

In the discussion below, we assume that the feasible domain of the LP is inscribed in the Euclidean unit ball (this can be ensured via standard scaling techniques). Notice this also implies that the feasible domain of the corresponding robust formulation is inscribed in the same ball.

A robust linear program with ellipsoidal noise is given by:

$$\begin{aligned} \exists? \quad & x \in \mathbb{R}^n \\ \text{s.t.} \quad & (a_i + Pu_i)^\top x - b_i \leq 0 \ , \quad \forall\ u_i \in \mathcal{U}, \quad i = 1, \ldots, m \ , \end{aligned} \tag{14}$$

where $P \in \mathbb{R}^{n \times K}$ is a matrix controlling the shape of the ellipsoidal uncertainty, $a_i \in \mathbb{R}^n$ are the *nominal* parameter vectors, and $\mathcal{U} = \{u \in \mathbb{R}^K \ : \ \|u\|_2 \leq 1\}$ is the K-dimensional Euclidean unit ball.

Dual-Subgradient Algorithm. The robust linear program (14) is amenable to our OGD-based meta-algorithm (Algorithm 1), as the constraints are linear with respect to the noise terms u_i. In this case we have $\nabla_u f_i(x, u) = P^\top x$, so that in each iteration of the algorithm, the update of the variables u_i^t takes the simple form

$$u_i^t \ \leftarrow \ \frac{u_i^{t-1} + \eta P^\top x}{\max\left\{\|u_i^{t-1} + \eta P^\top x\|_2, 1\right\}} \ .$$

Specializing Theorem 3 to the case of robust LPs, we obtain the following.

Corollary 10. *Algorithm 1 returns an ϵ-approximate robust solution to* (14) *after at most $T = O(\sigma^2/\epsilon^2)$ calls to the LP oracle, where $\sigma^2 = \|P\|_F^2$ is the maximal magnitude of the noise.*

Proof. Note that for all $u \in \mathcal{U}$ and $\|x\|_2 \leq 1$,

$$\|\nabla_u f_i(x, u)\|_2^2 \ = \ x^\top P P^\top x \ \leq \ \|P\|_F^2 \cdot \|x\|_2^2 \ \leq \ \sigma^2 \ .$$

Setting $D = 2$ and $G = \sigma$ in Theorem 3, we obtain the statement. □

Dual-Perturbation Algorithm. Since the constraints of the robust LP are linear in the uncertainties u_i, we can also apply our FPL-based meta-algorithm to the problem (14). Using the notations of Section 3.2, we have $\boldsymbol{g}_i(x) = P^\top x$. Hence, the computation of the noise variables u_i^t can be done in closed-form, as follows:

$$u_i^t \;\leftarrow\; \max_{\|u\|_2 \le 1} u^\top \big(P^\top \sum\nolimits_{\tau=1}^{t-1} x^\tau + p_t\big) \;=\; \frac{P^\top \sum_{\tau=1}^{t-1} x^\tau + p_t}{\|P^\top \sum_{\tau=1}^{t-1} x^\tau + p_t\|_2} \;.$$

In this case, Theorem 4 implies:

Corollary 11. *With high probability, Algorithm 2 returns an ϵ-approximate robust solution to (14) after at most $T = \tilde{O}(K\sigma^2/\epsilon^2)$ calls to the LP oracle, where $\sigma^2 = \|P\|_F^2$ is the maximal magnitude of the noise.*

Proof. Using the notations of Section 3.2 with $\boldsymbol{g}_i(x) = P^\top x$, we have

$$\begin{aligned} D &= \max_{u,v \in \mathcal{U}} \|u - v\|_1 \;\le\; 2\sqrt{K} \;, \\ G &= \max_{\|x\|_2 \le 1} \|P^\top x\|_1 \;\le\; \sqrt{K} \max_{\|x\|_2 \le 1} \|P^\top x\|_2 \;\le\; \sqrt{K}\,\sigma \;, \\ F &= \max_{\|x\|_2, \|u\|_2 \le 1} |x^\top P u| \;\le\; \|x\|_2 \cdot \|P\|_F \cdot \|u\|_2 \;\le\; \sigma \;. \end{aligned}$$

Making the substitutions into the guarantees in Theorem 4 completes the proof. □

We see that the asymptotic performance of Algorithm 2 is factor-K worse than that of Algorithm 1, in the case of robust LP problems.

4.2 Robust Quadratic Programming

A quadratically constrained quadratic program (QCQP) is given by

$$\begin{aligned} \exists? \quad & x \in \mathbb{R}^n \\ \text{s.t.} \quad & \|A_i x\|_2^2 - b_i^\top x - c_i \le 0 \;, \quad i = 1, \ldots, m \;, \end{aligned}$$

with $A_i \in \mathbb{R}^{n \times n}$, $b_i \in \mathbb{R}^n$, $c_i \in \mathbb{R}$. As in the case of LPs, we assume that the feasible domain of the above program is inscribed in the Euclidean unit ball. The robust counterpart of this optimization problem is a semidefinite program (Ben-Tal et al., 2009; Bertsimas et al., 2011). Current state-of-the-art QP solvers can handle two to three orders of magnitude larger QPs than SDPs, motivating our results. Indeed, our approach avoids reducing the robust program into an SDP and approximates it using a QP solver.

A robust QP with ellipsoidal uncertainties is given by[3]

$$\begin{aligned} \exists? \quad & x \in \mathbb{R}^n \\ \text{s.t.} \quad & \Big\|(A_i + \sum\nolimits_{k=1}^{K} u_{i,k} P_k)x\Big\|_2^2 - b_i^\top x - c_i \le 0 \;, \quad \forall\, u_i \in \mathcal{U}, \quad i = 1, \ldots, m \;, \end{aligned} \tag{15}$$

[3]For simplicity, the uncertainties we consider here are only in the matrices A_i (and not in the vectors b_i). In a similar, albeit more technical way we can also analyze our algorithm with general ellipsoidal uncertainties.

where $P_k \in \mathbb{R}^{n\times n}$ are fixed matrices and $\mathcal{U} = \{u \in \mathbb{R}^K \ : \ \|u\|_2 \le 1\}$. Here $u_{i,k}$ denotes the k'th entry of the noise vector $u_i \in \mathcal{U}$.

Notice that Algorithm 1 does not apply to formulation (15), as the constraints are certainly not concave with respect to the noise terms u_i (in fact, they are *convex* in u_i, as we show below). This motivates the need for our FPL-based meta-algorithm.

Dual-Perturbation Algorithm. We now show that the problem (15) falls in the scope of Section 3.2, and the assumptions required there hold for this program. We let $\sigma^2 = \sum_{k=1}^K \|P_k\|_F^2$ denote the total magnitude of the admissible noise, and assume that the Frobenius norms of the nominal matrices $A_1, \ldots, A_m$ are upper bounded by ρ.

The following lemma shows that the i'th constraint is in fact a convex quadratic in u_i.

Lemma 12. *The i'th constraint in* (15) *can be written as*

$$u_i^\top Q_x u_i + 2 r_x^\top u_i + s_x \ \le \ 0 \ ,$$

where $Q_x \in \mathbb{R}^{K\times K}$, $r_x \in \mathbb{R}^K$ and $s_x \in \mathbb{R}$ does not depend on u_i. The matrix Q_x is a positive semidefinite with $\|Q_x\|_F \le \sigma^2$, and $\|r_x\|_2 \le \sigma\rho$.

Proof. Define $y_0 = A_i x$ and $y_k = P_k x$ for $k = 1, \ldots, K$. We have

$$\begin{aligned}
\Big\|(A_i + \sum\nolimits_{k=1}^K u_{i,k} P_k)x\Big\|_2^2 \ &= \ x^\top \big(A_i^\top A_i + 2\sum\nolimits_{k=1}^K u_{i,k} A_i^\top P_k + \sum\nolimits_{k,l=1}^K u_{i,k} u_{i,l} P_k^\top P_l\big) x \\
&= \ y_0^\top y_0 + 2\sum\nolimits_{k=1}^K u_{i,k}\, y_0^\top y_k + \sum\nolimits_{k,l=1}^K u_{i,k} u_{i,l}\, y_k^\top y_l \\
&= \ y_0^\top y_0 + 2(Y^\top y_0)^\top u_i + u_i^\top Y^\top Y u_i \ ,
\end{aligned}$$

where $Y \in \mathbb{R}^{K\times K}$ is a matrix whose columns are $y_1, \ldots, y_K$. We see that the first claim holds for $Q_x = Y^\top Y$, $r_x = Y^\top y_0$ and $s_x = y_0^\top y_0 - b_i^\top x - c_i$, all of which are independent of u_i.

It is left to bound the coefficients in the above quadratic form. Note that $\|y_k\|_2 \le \|P_k\|_F \cdot \|x\|_2 \le \|P_k\|_F$ for all x with $\|x\|_2 \le 1$, so that

$$\|Y\|_F^2 \ = \ \sum_{k=1}^K \|y_k\|_2^2 \ \le \ \sum_{k=1}^K \|P_k\|_F^2 \ = \ \sigma^2 \ .$$

Hence,

$$\|Q_x\|_F \ = \ \|Y^\top Y\|_F \ \le \ \|Y\|_F^2 \ \le \ \sigma^2$$

and

$$\|r_x\|_2 \ = \ \|Y^\top y_0\|_2 \ \le \ \|Y\|_F \cdot \|y_0\|_2 \ \le \ \|Y\|_F \cdot \|A_i\|_F \ \le \ \sigma\rho \ ,$$

which proves the second claim. ∎

The above lemma demonstrates the well-known fact that the problem of finding the worst-case noise in robust QP with ellipsoidal uncertainty is a maximization of a convex quadratic over the unit ball (see Ben-Tal et al. 2009), a mathematical program known as the *trust region subproblem*. For this well-studied problem, fast approximation algorithms are available (Moré and Sorensen, 1983; Rendl and Wolkowicz, 1997) that are able to avoid solving an SDP (see also the recent linear-time approximation algorithm of Hazan and Koren 2014).

Finally, we compute the number of iterations required for Algorithm 2 to converge.

Corollary 13. *With high probability, Algorithm 2 returns an ϵ-approximate robust solution to (15) after at most $T = \tilde{O}(K^2\sigma^2\rho^2/\epsilon^2)$ calls to the QP oracle, where $\sigma^2 = \sum_{k=1}^{K} \|P_k\|_F^2$ is the total magnitude of the noise and $\rho \geq \max_i \|A_i\|_F$ is an upper bound over the norms of the nominal matrices.*

Proof. According to Lemma 12, the i'th constraint in (15) can be written in the linear form $\boldsymbol{g}_i(x) \cdot u_i' + h_i(x)$ by setting $\boldsymbol{g}_i(x) = (Q_x, 2r_x)$ and $h_i(x) = s_x$ and $u_i' = (u_i u_i^\top, u_i) \in \mathbb{R}^{K^2+K}$, with $Q_x \in \mathbb{R}^{K\times K}$, $r_x \in \mathbb{R}^K$, $s_x \in \mathbb{R}$ and $\|Q_x\|_F \leq \sigma^2$, and $\|r_x\|_2 \leq \sigma\rho$. That is, for the analysis only, we imagine that we work over a transformed, non-convex uncertainty set,

$$\mathcal{U}' \;=\; \left\{(uu^\top, u) \;:\; u \in \mathcal{U}\right\} \;\subseteq\; \mathbb{R}^{K^2+K} \;.$$

(Recall that the convergence properties of Algorithm 2 do not require the convexity of the uncertainty set.) Notice that maximizing the linear function $\boldsymbol{g}_i(x) \cdot u_i'$ with respect to $u_i' \in \mathcal{U}'$ is equivalent to maximizing the function $f_i(x, u_i)$ over $u_i \in \mathcal{U}$, as established by the oracle to the non-robust problem. With the above definitions and the notations of Section 3.2, we have for all $u' \in \mathcal{U}'$ that

$$\|u'\|_2 \;\leq\; \|uu^\top\|_F + \|u\|_2 \;\leq\; \|u\|_2^2 + \|u\|_2 \;\leq\; 2 \;,$$

and for all x, $\|x\|_2 \leq 1$ it holds that

$$\|\boldsymbol{g}_i(x)\|_2 \;\leq\; \|Q_x\|_F + 2\|r_x\|_2 \;\leq\; \sigma^2 + 2\sigma\rho \;.$$

Hence, for all u' and x,

$$\begin{aligned}
\|u'\|_1 \;&\leq\; \sqrt{K^2+K}\,\|u'\|_2 \;\leq\; 4K \;, \\
\|\boldsymbol{g}_i(x)\|_1 \;&\leq\; 2K\,\|\boldsymbol{g}_i(x)\|_2 \;\leq\; 2K(\sigma^2 + 2\sigma\rho) \;, \\
|\boldsymbol{g}_i(x) \cdot u'| \;&\leq\; \|\boldsymbol{g}_i(x)\|_2 \cdot \|u'\|_2 \;\leq\; 2\sigma^2 + 4\sigma\rho \;.
\end{aligned}$$

Hence, we may set

$$\begin{aligned}
D &= 8K \;, \\
G &= 2K(\sigma^2 + 2\sigma\rho) \;, \\
F &= 2\sigma^2 + 4\sigma\rho \;.
\end{aligned}$$

Theorem 4 with the above quantities now implies that Algorithm 2 needs at most $T = \tilde{O}(K^2\sigma^2\rho^2/\epsilon^2)$ iterations for ϵ-approximating the problem (15). □

4.3 Robust Semidefinite Programming

A semidefinite program (SDP) is given by

$$\begin{aligned} \exists? \quad & X \in \mathcal{S}_+^n \\ \text{s.t.} \quad & A_i \bullet X - b_i \leq 0 \ , \quad i = 1, \ldots, m \ , \end{aligned}$$

where $\mathcal{S}_+^n = \{X \in \mathbb{R}^{n \times n} \ : \ X \succeq 0\}$ is the cone of $n \times n$ positive semidefinite matrices, $A_i \in \mathbb{R}^{n \times n}$ are nominal parameter matrices, and $A \bullet X = \mathbf{Tr}(A^\top X)$ denotes the dot-product of the matrices A and X. Again, we assume that the feasible domain of the SDP is inscribed in the Euclidean unit ball (defined by the Frobenius matrix norm).

The robust counterpart of an SDP program is, in general, NP-hard even with simple ellipsoidal uncertainties (Ben-Tal and Nemirovski, 2002; Ben-Tal et al., 2009). Nevertheless, using our framework we are able to approximate robust SDP programs to within an arbitrary precision, as we now describe.

A robust SDP program with ellipsoidal uncertainties takes the following form:

$$\begin{aligned} \exists? \quad & X \in \mathcal{S}_+^n \\ \text{s.t.} \quad & \left(A_i + \sum\nolimits_{k=1}^{K} u_{i,k} P_k\right) \bullet X - b_i \leq 0 \ , \quad \forall \ u_i \in \mathcal{U}, \quad i = 1, \ldots, m \ , \end{aligned} \tag{16}$$

where $P_k \in \mathbb{R}^{n \times n}$ are fixed matrices and $\mathcal{U} = \{u \in \mathbb{R}^K \ : \ \|u\|_2 \leq 1\}$.

Dual-Subgradient Algorithm. Similarly to robust LPs, our OGD-based meta-algorithm can be applied to the robust SDP program (16) as the constraints are linear with respect to the noise terms u_i. In the present case we have $\nabla_u f_i(X, u) = (P_1 \bullet X, \ldots, P_K \bullet X)$, so that the update of the noise variables u_i^t takes the simple form

$$\forall \ k = 1, \ldots, K \ , \qquad v_{i,k}^t \ \leftarrow \ u_{i,k}^{t-1} + \eta (P_k \bullet X) \ , \qquad u_{i,k}^t \ \leftarrow \ \frac{v_{i,k}^t}{\max\{\|v_i\|_2, 1\}} \ .$$

For the resulting algorithm, we have the following.

Corollary 14. *Algorithm 1 terminates with an ϵ-approximate solution to* (16) *after no more than $T = O(\sigma^2/\epsilon^2)$ calls to the SDP oracle, where $\sigma^2 = \sum_{k=1}^K \|P_k\|_F^2$ is the total magnitude of the allowed noise.*

Proof. By Cauchy-Schwarz, for all $u \in \mathcal{U}$ and $\|x\|_2 \leq 1$,

$$\|\nabla_u f_i(x, u)\|_2^2 \ = \ \sum_{k=1}^{K} |P_k \bullet X|^2 \ \leq \ \sum_{k=1}^{K} \|P_k\|_F^2 \cdot \|X\|_F^2 \ \leq \ \sum_{k=1}^{K} \|P_k\|_F^2 \ = \ \sigma^2 \ .$$

Therefore, we may take $D = 2$ and $G = \sigma$ in the bound of Theorem 3, giving our claim. □

Finally, we note that Algorithm 2 also applies to robust SDPs, but gives a guarantee worse by a factor of K^2.

5 Conclusion

In this paper we considered using online learning approaches for effectively solving robust optimization problems without transforming the problem to a different, more complex, class of problems. We showed that if the original problem is convex and comes equipped with an oracle capable of approximating it, then we can solve the robust problem approximately by employing an online learning approach that invokes the oracle a polynomial number of times. Essentially, our approach is applicable to any robust optimization problem for which we can efficiently approximate the worst-case noise for any given feasible solution, and is particularly efficient when the latter task can be accomplished via subgradient methods.

Our approach opens up avenues for solving large-scale robust optimization problems that are more common in data analysis and machine learning. The key observation is that the number of iterations of the online learning algorithms is *independent* of the dimension of the problem. This means that as long as the original problem is solvable efficiently (e.g., support vector machines) the robust problem does not become much more difficult if the accuracy of the solution can be compromised.

Our approach can be used to solve other RO problems of interest. For example, solving robust multi-stage decision problems such as Markov decision processes (Puterman, 1994) is of interest; see Nilim and El Ghaoui (2005) for discussion of robust Markov decision processes. Standard (non-robust) Markov decision processes are solvable using linear programming. However, their robust counterpart is in general not amenable to linear programming and a dynamic programming approach is needed to solve the stochastic game between the decision maker and Nature. This approach does not seem to scale up to large problems where approximate dynamic programming is needed. Using an online approach as we suggested may prove very useful since solving the original problem seems easy (solving a linear program) and finding the worst-case noise is also not too difficult, depending on the noise model. We leave the important case of multi-stage problems for future research.

Finally, it would be interesting to adapt our approach to robust combinatorial optimization, where few disciplined robust optimization methods are available. While our methods assume the original problem to be convex, our main interaction with the problem is through a black-box oracle (that may be available for non-convex problems), so it seems that the convexity requirement might be relaxed in certain cases of interest.

References

S. Arora, E. Hazan, and S. Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing*, 8(6):121–164, 2012.

A. Ben-Tal and A. Nemirovski. Robust optimization - methodology and applications. *Math. Program.*, 92(3):453–480, 2002.

A. Ben-Tal, L. E. Ghaoui, and A. Nemirovski. *Robust Optimization*. Princeton Series in Applied Mathematics. Princeton University Press, October 2009.

D. Bertsimas and M. Sim. Robust discrete optimization and network flows. *Math. Program.*, 98(1-3):49–71, 2003.

D. Bertsimas, D. B. Brown, and C. Caramanis. Theory and applications of robust optimization. *SIAM Review*, 53(3):464–501, 2011.

C. Bhattacharyya, L. R. Grate, M. I. Jordan, L. El Ghaoui, and I. S. Mian. Robust sparse hyperplane classifiers: Application to uncertain molecular profiling data. *Journal of Computational Biology*, 11(6):1073–1089, 2004a.

C. Bhattacharyya, K. S. Pannagadatta, and A. J. Smola. A second order cone programming formulation for classifying missing data. In L. K. Saul, Y. Weiss, and L. Bottou, editors, *Advances in Neural Information Processing Systems (NIPS17)*, Cambridge, MA, 2004b. MIT Press.

G. Calafiore and M. C. Campi. Uncertain convex programs: randomized solutions and confidence levels. *Mathematical Programming*, 102(1):25–46, 2004.

N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, New York, 2006.

K. L. Clarkson, E. Hazan, and D. P. Woodruff. Sublinear optimization for machine learning. *J. ACM*, 59(5):23:1–23:49, 2012.

D. Garber and E. Hazan. Approximating semidefinite programs in sublinear time. In *25th Annual Conference on Neural Information Processing Systems (NIPS)*, pages 1080–1088, 2011.

E. Hazan. The convex optimization approach to regret minimization. *Optimization for machine learning*, 1, 2011.

E. Hazan and T. Koren. A linear-time algorithm for trust region problems. *arXiv preprint arXiv:1401.6757*, 2014.

E. Hazan, T. Koren, and N. Srebro. Beating sgd: Learning svms in sublinear time. In *Advances in Neural Information Processing Systems*, pages 1233–1241, 2011.

A. T. Kalai and S. Vempala. Efficient algorithms for online decision problems. *J. Comput. Syst. Sci.*, 71(3):291–307, 2005.

G. R. Lanckriet, L. El Ghaoui, C. Bhattacharyya, and M. I. Jordan. A robust minimax approach to classification. *Journal of Machine Learning Research*, 3:555–582, 2003.

J. J. Moré and D. C. Sorensen. Computing a trust region step. *SIAM Journal on Scientific and Statistical Computing*, 4(3):553–572, 1983.

A. Mutapcic and S. P. Boyd. Cutting-set methods for robust convex optimization with pessimizing oracles. *Optimization Methods and Software*, 24(3):381–406, 2009.

A. Nilim and L. El Ghaoui. Robust control of Markov decision processes with uncertain transition matrices. *Operations Research*, 53(5):780–798, September 2005.

S. A. Plotkin, D. B. Shmoys, and É. Tardos. Fast approximation algorithms for fractional packing and covering problems. *Mathematics of Operations Research*, 20(2):257–301, 1995.

M. L. Puterman. *Markov Decision Processes*. John Wiley & Sons, New York, 1994.

F. Rendl and H. Wolkowicz. A semidefinite framework for trust region subproblems with applications to large scale minimization. *Mathematical Programming*, 77(1):273–299, 1997.

S. Shalev-Shwartz. Online learning and online convex optimization. *Found. Trends Mach. Learn.*, 4(2):107–194, 2012.

P. K. Shivaswamy, C. Bhattacharyya, and A. J. Smola. Second order cone programming approaches for handling missing and uncertain data. *Journal of Machine Learning Research*, 7:1283–1314, July 2006.

T. Trafalis and R. Gilbert. Robust support vector machines for classification and computational issues. *Optimization Methods and Software*, 22(1):187–198, February 2007.

H. Xu, C. Caramanis, and S. Mannor. Robustness and regularization of support vector machines. *Journal of Machine Learning Research*, 10(Jul):1485–1510, 2009.

H. Xu, C. Caramanis, and S. Mannor. Robust regression and lasso. *IEEE Transactions on Information Theory*, 56(7):3561 – 3574, 2010.

M. Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *ICML*, pages 928–936, 2003.